Algebra Universalis

Check for
updates

Saturated free algebras and almost indiscernible theories

Thomas G. Kucera and Anand Pillay

Abstract. We extend the concept of "almost indiscernible theory" introduced by Pillay and Sklinos in 2015 (which was itself a modernization and expansion of Baldwin and Shelah from 1983), to uncountable languages and uncountable parameter sequences. Roughly speaking a theory T is almost indiscernible if some saturated model is in the algebraic closure of an indiscernible set of sequences. We show that such a theory T is nonmultidimensional, superstable, and stable in all cardinals $\geq |T|$. We prove a structure theorem for sufficiently large a-models M, which states that over a suitable base, M is in the *algebraic closure* of an independent set of realizations of weight one types (in possibly infinitely many variables). We also explore further the saturated free algebras of Baldwin and Shelah in both the countable and uncountable context. We study in particular theories and varieties of R-modules, characterizing those rings R for which the free R-module on $|R|^+$ generators is saturated, and pointing out a counterexample to a conjecture by Pillay and Sklinos.

Mathematics Subject Classification. 08B20, 03C45, 03C05, 03C60, 16D40.

Keywords. Free algebras, Indiscernible sets, Forking independence, Free modules.

1. Preliminaries

This paper continues and builds on the investigations of Baldwin and Shelah [2] and Pillay and Sklinos [11]. The original context of Baldwin and Shelah was the study of $\mathrm{Th}(F)$ where F is the free algebra in $\aleph_1$-many generators in a variety (in the sense of universal algebra) in a countable signature or language. Their work was clarified (with some corrections) by Pillay and Sklinos, and also extended to the more general notion of *almost indiscernible theories*, still in a countable language.

Presented by M. Pinsker.

𝄐 Birkhäuser

As stated in the abstract, roughly speaking a theory T is almost indiscernible if some saturated model is in the algebraic closure of an indiscernible set of sequences. In the same casual manner, one might say that a free algebra is one that is freely generated by a basis. "Algebraic closure of ..." is a general model-theoretic analogue of "generated by ..."; "indiscernible sequence" is an analogue of "free (basis)". These analogies led Pillay and Sklinos to the generalizations of Baldwin-Shelah presented in [11].

Similarly, the present abstract generalization of those results was in part motivated by our understanding of the structure of nicely behaved theories of modules: we knew that something like Theorem 2.17 should be true, given the correct definitions and proper development of the theory.

Although some difficult technical results in classical stability theory in Section 2 provide the foundation for the entire paper, the ultimate goal is to pursue the applications to algebraic structures of various kinds.

Thus we have several aims in the current paper. First we consider the case of almost indiscernible theories but generalized to uncountable languages as well as indiscernible sets of infinite (rather than finite) tuples. Among interesting differences with the countable case is that the theories will be superstable but not necessarily totally transcendental. The main structural result is Theorem 2.10. We point out that the almost indiscernible (complete) theories T of modules (over some ring R in the usual language) are precisely the superstable theories of modules which are $|T|$-stable. We also revisit the special case of saturated free algebras, with respect to a given variety in the sense of universal algebra. Results from Pillay–Sklinos in the context of countable languages go through smoothly for uncountable languages. On the other hand in the even more special case of the variety of R-modules, if the free algebra is saturated then its theory is totally transcendental (not just superstable). We classify the rings R such that the free R-modules are saturated, and give an example where the corresponding theory does not have finite Morley rank, yielding a counterexample to a question from Pillay–Sklinos.

We only use very standard facts from Shelah's stability theory. Rather than tracing back to original sources, we rely primarily on the outline in Chapter 1 of Pillay [8], with occasional reference to Baldwin [1] where required.

In our work, we follow the proofs of [11] as closely as possible, but give some further clarifications to the structure of the proofs of Propositions 2.9 and 2.10 therein, as necessitated by the move to the uncountable context.

1.1. The definition

We begin by extending Definition 2.1 of Pillay and Sklinos [11].

Definition 1.1. Let T be a complete theory of cardinality τ, $\mu \leq \tau$ a finite or infinite cardinal, and $\kappa > \tau$ a cardinal.

T is called (μ, κ)-*almost indiscernible* if it has a saturated model M of cardinality κ which is in the algebraic closure of an indiscernible set I of μ-sequences.

T is *almost indiscernible* if it is (μ, κ)-almost indiscernible for some such μ, κ.

So "almost indiscernible" as defined in [11] for countable theories is $(n, \aleph_1)$-almost indiscernible for some finite $n > 0$.

Trivially if T is (μ, κ)-almost indiscernible then it is (μ', κ)-almost indiscernible for any μ', $\mu \leq \mu' \leq \tau$, for given an indiscernible set of μ-sequences, just extend each to a μ'-sequence by repeating the first entry. As well, in the definition, we can replace sequences indexed by μ by sequences of cardinality μ.

Example 1.2. Later (Corollary 2.3) we will see that if T is almost indiscernible, it even has a saturated model of cardinality τ which is the algebraic closure of an indiscernible set of μ-tuples.

But we should not include the possibility of $\kappa = \tau$ in the definition. Let $\langle \mathcal{Q}; \leq \rangle$ be the disjoint union of countably many copies of the rational linear order, let $\mathbf{e}$ be some fixed enumeration of the rationals $\mathbb{Q}$ in order type ω, and let $\underline{e}_i$ be the copy of this tuple on the i-th copy of $\mathbb{Q}$ in $\mathcal{Q}$. Then clearly $\langle \underline{e}_i : i < \omega \rangle$ is an indiscernible set of ω-tuples in $\mathcal{Q}$ whose algebraic closure (in fact union) is all of $\mathcal{Q}$.

Of course the theory of this structure cannot have any of the other properties of an almost indiscernible theory, expounded later.

Example 1.3. The definition does not require "best possible choices". Let $\mathbb{F}$ be an infinite field of cardinality τ, $\mathcal{L}$ the usual language for vector spaces over $\mathbb{F}$, and T the theory of non-zero vector spaces over $\mathbb{F}$. Let $(\underline{e}_\alpha)_{\alpha < \tau^+}$ enumerate a basis of the $\mathbb{F}$-vector space $\mathcal{V}$ of dimension τ^+. Then clearly $\{ \underline{e}_\alpha : \alpha < \tau^+ \}$ is an indiscernible set generating $\mathcal{V}$ and so T is $(1, \tau^+)$-almost indiscernible.

But for any cardinal $1 < \mu \leq \tau$, we could just as well take $\underline{e}_0$ to be a μ-tuple enumerating a basis of the $\mathbb{F}$-vector space of dimension μ, and extend it to a sequence $(\underline{e}_\alpha)_{\alpha < \tau^+}$ whose range is again a basis of $\mathcal{V}$, exhibiting T as a (μ, τ^+)-almost indiscernible theory.

Alternatively, we could take $\underline{e}_0$ to be an enumeration of $\mathbb{F}$, that is, of the one-dimensional vector space, and then let $(\underline{e}_\alpha)_{\alpha < \tau^+}$ be an enumeration of τ^+ direct-sum independent subspaces of $\mathcal{V}$, exhibiting T as a (τ, τ^+)-almost indiscernible theory, with a lot more information than is really required.

One of the goals of the structure theory we develop is to recover some of the fine detail that might be lost by redundancy in the indiscernible set of I-sequences.

Example 1.4. We point out briefly that both the condition of almost indiscerniblity, as well as the consequences in Theorem 2.10 below, concern exceptional behaviour, even for uncountably categorical theories T. Of course when T is almost strongly minimal, that is, any model is in the algebraic closure of a fixed strongly minimal set D (without parameters say), then we do have almost indiscernibility: any model M is in the algebraic closure of the indiscernible set consisting of a maximal independent set of realizations in M of the generic type of D. Even when T is not almost strongly minimal, such as $\mathrm{Th}(\mathbb{Z}/4\mathbb{Z})^{(\omega)}$, the conditions of almost indiscernibility may still hold. Let 1 be the generator of a copy of $\mathbb{Z}/4\mathbb{Z}$: then $\mathrm{tp}(1)$ has Morley rank 2 but weight 1, and of course any model is generated as a $\mathbb{Z}/4\mathbb{Z}$-module by an independent set of realizations of this type.

However consider the theory T of the structure consisting of two sorts X and V with surjective $\pi : X \to V$, and where V has the structure of an infinite-dimensional vector space over $\mathbb{F}_2$ say, and each fibre X_a is (uniformly definably) a principal homogeneous space for $(V, +)$. Clearly T can be axiomatized in a two-sorted language with symbols π of sort $X \to V$, $+$ of sort $V \times V \to V$, and $\langle \, , \rangle$ of sort $V \times X \to X$. Then T is $\aleph_1$-categorical, but is *not* almost indiscernible. V is strongly minimal and X has Morley rank 2, degree 1. Let a be a generic point of V and $b \in X$ be such that $\pi(b) = a$. Then $\mathrm{tp}(b)$ has weight 1. If $\{\, b_i : i \in I \,\}$ is a maximal independent set of realizations of $\mathrm{tp}(b)$ in a model M, then necessarily $\{\, \pi(b_i) : i \in I \,\}$ is a linearly independent set in V. Furthermore, $\mathrm{acl}(V) = V$, so for distinct non-zero $u, v \in V$, no part of the fibre over $u + v$ is algebraic over the (union of the) fibres over u and v. So on the one hand any maximal indiscernible set I (of tuples) in a model M cannot intersect all the fibres of π, and on the other hand any fibre that does not intersect I is not in the algebraic closure of I.

1.2. Basic facts

We need some translations of the basic facts about stability theory enunciated in Section 1 of [11]. Our theories will be superstable, not necessarily totally transcendental, and so may not have prime models. As is usual, we will abbreviate "superstable" as "ss" and "totally transcendental" as "tt". We remind the reader that for *uncountable* languages we have to characterize tt theories as those where every formula (every type) has ordinal-valued Morley rank. For countable theories only, this is equivalent to ω-stability. But each theory T that we consider will nonetheless be stable in $\tau = |T|$, and will have a model M_ω which is ω-saturated, is an a-model, and such that every stationary type is non-orthogonal to a type over M_ω. Furthermore, *every* model we care about is an elementary extension of M_ω, so there are still very strong parallels to [11]. The principal difficulty lies not so much in the movement to merely superstable theories, but in allowing infinite sequences as the elements of the indiscernible sets.

For the remainder of this section and in Section 2 (unless explicitly stated otherwise), T is a **complete superstable theory** in a language of cardinality τ and $\mathfrak{M}$ is a sufficiently large saturated model of T (a universe), with $\mathfrak{M}^{\mathrm{eq}}$ being the associated "imaginary" universe. We work in $\mathfrak{M}^{\mathrm{eq}}$: every element, set, sequence, model that we consider is a "small" thing in $\mathfrak{M}^{\mathrm{eq}}$, that is, of cardinality strictly less than the cardinality of $\mathfrak{M}$. By "algebraic (or definable) closure" we always mean "in the sense of $\mathfrak{M}^{\mathrm{eq}}$". However, in reading the algebraic examples, it is always helpful to think of things taking place in the "home" sort.

Fact 1.5 [8, Chapter 1, Lemmas 4.1.1, 4.1.2, 4.2.1]. There is a cardinal $\lambda(T) \leq 2^\tau$ such that T is stable in κ iff $\kappa \geq \lambda(T)$. T has a saturated model in every cardinal $\kappa \geq \lambda(T)$. Since T is superstable, $\kappa(T) = \aleph_0$, that is, every type (in finitely many variables) does not fork over some finite set.

Recall that if A, B, C are sets of parameters (or tuples), then B *dominates* C *over* A if whenever D is independent from B over A then D is independent from C over A.

Definition 1.6. (a) The *strong type* of a over A, $\mathrm{stp}(a/A)$, is the type of a over $\mathrm{acl}(A)$ (for emphasis, in $\mathfrak{M}^{\mathrm{eq}}$).

  (b) By an *a-model* of T we mean a model M of T such that any strong type over any finite subset of M is realized in M.

  (c) A type $p(x) \in S(A)$ is said to be *a-isolated* if there is a finite subset B of A and a strong type $q(x)$ over B which implies $p(x)$.

Fact 1.7 [8, Chapter 1, Lemmas 4.2.4, 4.3.4].

  (a) For any set of parameters A there is an *a*-prime model over A, that is, an *a*-model M containing A such that for any *a*-model N containing A there is an elementary embedding over A of M into N. M has the property that for all tuples b from M, $\mathrm{tp}(b/A)$ is *a*-isolated.

  (b) Suppose M_0 is an *a*-model, and A is any set of parameters and b any tuple. Then $\mathrm{tp}(b/M_0A)$ is *a*-isolated iff A dominates Ab over M_0.

Clearly an *a*-model is $\aleph_0$-saturated.

Definition 1.8. T is *nonmultidimensional* ["nmd"] if every stationary type p is nonorthogonal to $\emptyset$, that is, nonorthogonal to some stationary type which does not fork over $\emptyset$.

Remark 1.9. Definition 1.8 is equivalent to the following:

For any A and any stationary type $q(x)$ over A, if $\mathrm{stp}(A'/\emptyset) = \mathrm{stp}(A/\emptyset)$, A' is independent from A over $\emptyset$, and q' is the copy of q over A', then q is nonorthogonal to q'.

In fact, at least for superstable theories, in 1.8 it suffices to demand that every stationary type p be non-orthogonal to a type over a fixed *a*-model (as mentioned in the first paragraph of this Section), cf. Baldwin [1, XV: Theorem 1.7].

We need the following basic result:

Proposition 1.10. *Let T be a superstable nonmultidimensional theory (of any cardinality). Then any elementary extension of an a-model is an a-model.*

This proposition is folklore and there are various routes to it. For example it follows directly from Shelah's 'three model lemma', and also follows from Propositions 3.2 and 3.6 of Chapter 7 of [8]. We will give a quick independent proof, starting with a suitable 3-model lemma.

Lemma 1.11. *Assume that T is superstable nonmultidimensional, and that $M_0 \preceq M \prec N$ where M_0 is an a-model. Then there is $c \in N \setminus M$ such that $\mathrm{tp}(c/M)$ is regular and does not fork over M_0.*

Proof. Choose $b \in N \backslash M$ such that $\mathrm{R}^{\infty}(\mathrm{tp}(b/M)) = \alpha$ is minimized. Let $\varphi(x, a)$ with $a \in M$ be a formula in $\mathrm{tp}(b/M)$ of ∞-rank α. In particular $\mathrm{tp}(b/M)$ does not fork over a. As M_0 is an a-model we can choose $a' \in M_0$ such that $\mathrm{stp}(a') = \mathrm{stp}(a)$ and a' is independent from a over $\emptyset$. By Remark 1.9, there is a type $q(x)$ over M_0 which contains the formula $\varphi(x, a')$ and is nonorthogonal to $\mathrm{tp}(b/M)$. So there is $M' \supseteq M$ with b independent from M' over M and c' realizing $q|M'$ such that b forks with c' over M'.

A standard argument yields $c \in N \backslash M$ such that $\models \varphi(c, a')$: There is a formula $\chi(x, y, z)$ over M and $d \in M'$ such that $\models \chi(b, c', d)$ witnesses the forking of b with c' over M', that is, $\chi(x, c'', d'') \cup \mathrm{tp}(b/M)$ forks over M for any c'', d''. Now $\exists y(\chi(x, y, d) \wedge \varphi(x, a'))$ is in $\mathrm{tp}(b/M')$ so for some $d' \in M$, $\exists y(\chi(x, y, d') \wedge \varphi(x, a'))$ is in $\mathrm{tp}(b/M)$. Let $c \in N$ be such that $\models \chi(b, c, d') \wedge \varphi(c, a')$. As b forks with c over M, $c \in N \backslash M$.

So as $\mathrm{R}^{\infty}(\varphi(x, a')) = \alpha$, by the minimal choice of α, $\mathrm{R}^{\infty}(\mathrm{tp}(c/M) = \alpha$. Then as $a' \in M_0$, $\mathrm{tp}(c/M)$ does not fork over M_0, as required. Furthermore, $\mathrm{tp}(c/M)$ is regular by the choice of α minimal, as in the argument of [8, Lemma 4.5.6]. □

Proof (of Proposition 1.10). T is assumed to be superstable nonmultidimensional. Let M be an a-model, and assume $M \prec N$. We want to prove that N is an a-model. Let N' be the a-prime model over N given by Fact 1.7(a). It will be enough to show that $N = N'$.

Suppose not. Then by Lemma 1.11, there is $c \in N' \backslash N$ which is independent from N over M. But $\mathrm{tp}(c/N)$ is a-isolated, so by Fact 1.7(b), N dominates c over M which is a contradiction. □

Corollary 1.12 (*T superstable nonmultidimensional*) *If M is an a-model and N is a-prime over $M \cup A$, then N is prime and minimal over $M \cup A$.*

Proof. It is immediate by 1.10 that N is prime; and by the proof just given, it is minimal over $M \cup A$. □

Now let $\mathcal{A}_0$ be the a-prime model of T over $\emptyset$.

Let $(p_i)_{i \in I}$ be a list, up to non-orthogonality, of all the regular types over $\mathcal{A}_0$ (and hence, up to non-orthogonality, all the regular types of T). Since T is superstable, for each $i \in I$ there is finite $a_i \in \mathcal{A}_0$ such that p_i is definable over a_i. [Since T is superstable, we can choose finite $b \in \mathcal{A}_0$ so that p_i does not fork over b; then since $\mathcal{A}_0$ is an a-model, we can find $c \in \mathcal{A}_0$ realizing the restriction of p_i to b and in the correct strong type. Then p_i is definable over $a_i = bc$.] We let $\hat{p}_i$ be the restriction of p_i to $\mathrm{acl}(a_i)$.

Comment. Note that in general $|I| \leq 2^{|T|}$, but we will see that in the context we develop, as T has an a-prime model of cardinality $\tau = |T|$ and T is τ-stable, in fact $|I| \leq \tau$.

We need a slight reformulation of [11, Fact 1.3].

Lemma 1.13. *Let $\mathcal{A}_0 \prec M \prec \mathfrak{M}$. For each $i \in I$ let J_i be a maximal independent set of realizations of p_i in M. Then M is a-prime, prime, and minimal over $\mathcal{A}_0 \cup \bigcup_{i \in I} J_i$.*

2. Almost indiscernible theories

Context 2.1. Unless explicitly stated otherwise, for Section 2:

(a) T is a (μ, κ)-*almost indiscernible* theory, $|T| = \tau$, $\mu \leq \tau < \kappa$, with universe $\mathfrak{M}$ of some regular cardinality $\overline{\kappa} \gg \kappa$.

(b) M is a saturated model as in the definition: $|M| = \kappa$, I is an indiscernible *set* of μ-sequences in M, and M is in the algebraic closure of the (union of) I.

Since $\mu \leq \tau < \kappa$, necessarily $|I| = \kappa$, so we can write I as a κ-sequence $\langle \underline{e}_\alpha : \alpha < \kappa \rangle$, and when necessary the μ-sequence $\underline{e}_\alpha$ is indexed as $\langle \underline{e}_{\alpha,i} : i < \mu \rangle$.

(c) Extend I to an indiscernible 'set' $\overline{I} = \langle \underline{e}_\alpha : \alpha < \overline{\kappa} \rangle$ in $\mathfrak{M}$.

For each infinite *ordinal* $\lambda \leq \overline{\kappa}$, let $I_\lambda = \langle \underline{e}_\alpha : \alpha < \lambda \rangle$ and set $M_\lambda = \mathrm{acl}(I_\lambda)$ in $\mathfrak{M}$. Note the extension of this sequence to the size of the universe.

In particular, $M_\kappa = M$ is an elementary substructure of $\mathfrak{M}$, but the status of all the other M_λ remains to be resolved.

2.1. Basic facts

Recall that for an infinite cardinal ν, M is F^a_ν-saturated if every strong type over any subset of M of cardinality less than ν is realized in M.

Theorem 2.2. λ *denotes an infinite* ordinal.

(a) $\lambda \geq \tau$ *implies* $|M_\lambda| = |\lambda|$ *and* $\lambda < \tau$ *implies* $\mu\,|\lambda| \leq |M_\lambda| \leq \tau$.

(b) *For all* $\lambda \leq \overline{\kappa}$, $M_\lambda \preceq \mathfrak{M}$.

(*And so* $\langle M_\lambda \rangle_{\omega \leq \lambda \leq \overline{\kappa}}$ *is an elementary chain.*)

(c) *For all* $\lambda \leq \overline{\kappa}$, M_λ *is* $|\lambda|$-*saturated.*

(d) *In particular* $M_{\overline{\kappa}}$ *is saturated of cardinality* $\overline{\kappa}$, *so is without loss of generality equal to* $\mathfrak{M}$.

(e) *For all* $\lambda \leq \overline{\kappa}$, M_λ *is* $F^a_{|\lambda|}$-*saturated.*

(f) *In particular all* M_λ *have the property that all strong types over finite sets are realized.*

Then it will follow from Theorem 2.4 *to come, that since* T *is superstable, all the* M_λ *are a-models.*

Proof.   (a) follows by simple counting.

(b) By definition $M = M_\kappa$ is an elementary substructure of $\mathfrak{M}$. First assume $\lambda < \kappa$. Then an easy Tarski-Vaught argument which we now describe yields that $M_\lambda \preceq M_\kappa$: Let $\varphi(x)$ be a formula with parameters from M_λ with a solution $d \in M_\kappa$. Let J be a finite subset of λ and $\alpha_1, \ldots, \alpha_n$ distinct elements of $\kappa \backslash \lambda$, and such that the parameters in $\varphi(x)$ are from $E = \mathrm{acl}(\{\,\underline{e}_\beta : \beta \in J\,\})$ and $d \in \mathrm{acl}(E \cup \{\,\underline{e}_{\alpha_i} : i = 1, \ldots, n\,\})$. Let $\gamma_1, \ldots, \gamma_n$ be distinct elements of $\lambda \backslash J$. Then by indiscernibility, $\mathrm{tp}(\underline{e}_{\gamma_1}, \ldots, \underline{e}_{\gamma_n}/E) = \mathrm{tp}(\underline{e}_{\alpha_1}, \ldots, \underline{e}_{\alpha_n}/E)$ (where these types are computed in M_κ). Thus we can find $d' \in \mathrm{acl}(E, \underline{e}_{\gamma_1}, \ldots, \underline{e}_{\gamma_n}) \subseteq M_\lambda$ such that $M_\kappa \models \varphi(d')$.

On the other hand, assume that $\lambda > \kappa$. We will also use Tarski-Vaught: Let $\varphi(x)$ be a formula over M_λ such that $\mathfrak{M} \models \exists x \varphi(x)$. Let J be a finite subset

of λ such that the parameters of φ are in $E = \mathrm{acl}(\{\,\underline{e}_\beta : \beta \in J\,\})$. Choose a subset J' of κ of the same cardinality as J, so by indiscernibility J and J' have the same type (under any enumeration) in $\mathfrak{M}$. Let $E' = \mathrm{acl}(J')$, and let $\varphi'(x)$ be the image of $\varphi(x)$ under a partial elementary map f taking J to J' and E to E'. As $M_\kappa \prec \mathfrak{M}$, $M_\kappa \models \exists x \varphi'(x)$. So let $d' \in M_\kappa$ be such that $M_\kappa \models \varphi(d')$. Again d' will be in $\mathrm{acl}(J', K')$ where K' is a finite subset of κ disjoint from J'. Let K be a finite subset of λ of the same cardinality as K' and disjoint from J. Then again indiscernibility implies that the partial elementary map f (in the sense of $\mathfrak{M}$) extends to a partial elementary map g taking K to K' and $\mathrm{acl}(E, K)$ to $\mathrm{acl}(E', K')$. Let $g(d) = d'$, so $\mathfrak{M} \models \varphi(d)$ and $d \in M_\lambda$.

(c) Note that this part is an analogue of Morley's theorem (that for a countable complete theory T, if for some uncountable cardinal κ all models of T of cardinality κ are saturated, then all uncountable models of T are saturated), and our proof of the harder case (where $\lambda > \kappa$) will be closely related to the proof of Morley's theorem as given in [7, Theorem 5.33].

The first (and easy) case is when $\lambda < \kappa$. Then the proof of part (b) (in the case $\lambda < \kappa$) adapts. Namely in this case we have a type $p(x)$ over a set of parameters $E \subset M_\lambda$ of cardinality $< |\lambda|$, and now choose $J \subset \lambda$ of cardinality $< |\lambda|$ with $E \subseteq \mathrm{acl}(J)$ (without loss $E = \mathrm{acl}(J)$). p is realized in M_κ by some d in the algebraic closure of E together with finitely many $\underline{e}_\alpha$ with $\alpha \in \kappa \setminus J$. Then as $|J| < |\lambda|$, we can again replace these $\underline{e}_\alpha$ by some $\underline{e}_\gamma$ in M_λ and realize p in M_λ.

The harder case occurs when $\lambda > \kappa$.

There is no harm (for notational simplicity) in assuming that λ is a cardinal. We suppose that M_λ is not saturated, and aim for a contradiction. Then there is a subset A of M_λ with $|A| < \lambda$ and a complete type $p(x)$ over A which is not realized in M_λ. We may assume that A has cardinality at least τ. Let $J \subseteq I_\lambda$ be of cardinality $|A|$ such that $A \subseteq \mathrm{acl}(J)$, and let I be a countable subset of I_λ disjoint from J. Extending p to a complete type over $\mathrm{acl}(J)$ we may assume that $A = \mathrm{acl}(J)$. Note that I is indiscernible over A. Note also that, by part (b) $\mathrm{acl}(A, I)$ is an elementary substructure N of $\mathfrak{M}$ (and of M_λ), and $p(x)$ is not realized in N. In particular for any consistent formula $\varphi(x)$ over $A \cup I$, we can pick a formula $\psi_\varphi(x) \in p(x)$ such that $\models \exists x (\varphi(x) \wedge \neg \psi_\varphi(x))$. (Here "$\models$" means, equivalently, in $\mathfrak{M}$ or in M_λ or in N.)

We now construct a subset J' of J of cardinality τ, and $A' = \mathrm{acl}(J') \subseteq A$ such that for $p'(x) = p|A'$ we have:

(*) For each consistent formula $\varphi(x)$ over $A' \cup I$, there is a formula $\psi(x) \in p'(x)$ such that $\models \exists x (\varphi(x) \wedge \neg \psi(x))$.

We do this by a routine union of chain argument. We define a sequence of pairs $\langle J_i, A_i \rangle$, $J_i \subseteq J$, $|J_i| = \tau$, and $A_i = \mathrm{acl}(J_i) \subseteq A$ by recursion on $n < \omega$.

Let J_0 be any subset of J of cardinality τ, and $A_0 = \mathrm{acl}(J_0) \subseteq A$.

Given $\langle J_i, A_i \rangle$, for each consistent formula $\varphi(x)$ over $A_i \cup I$, we have $\psi_\varphi(x) \in p(x)$ such that $\models \exists x (\varphi(x) \wedge \neg \psi_\varphi(x))$. Note that there are at most τ such formulas $\psi_\varphi(x)$. Add the parameters (from A) of the formulas $\psi_\varphi(x)$ to A_i to obtain A_i' (which still has cardinality τ). Extend J_i to $J_{i+1} \subseteq J$ of cardinality τ such that $A_i' \subseteq \mathrm{acl}(J_{i+1})$, and set $A_{i+1} = \mathrm{acl}(J_{i+1})$. Set $A' = \bigcup_n A_n$ and

$J' = \bigcup_n J_n$. Then set $p'(x)$ to be the restriction of p to A'. So we have obtained (*).

Now let I' be a subset of I_λ of cardinality κ which is disjoint from J'. Then $\mathrm{acl}(J', I')$ is an elementary substructure M' of $\mathfrak{M}$ which is isomorphic to M_κ (as $|J' \cup I'| = \kappa$). By hypothesis M' is κ-saturated, and A' has cardinality $\tau < \kappa$, so the type $p'(x)$ is realized in M', by some d'. But $d' \in \mathrm{acl}(A', I')$ so its type over $A' \cup I'$ is isolated by a consistent algebraic formula $\varphi(x)$ over $A' \cup I'$. We will exhibit the parameters from I' by writing φ as $\varphi(x, b_1, \ldots, b_n)$ where b_i is a finite tuple from some $\underline{e}_i \in I'$. Now $M \models \forall x(\varphi(x, b_1, .., b_n) \to \psi(x))$ for all $\psi(x) \in p'(x)$, as $\varphi(x, b_1, \ldots, b_n)$ isolates $\mathrm{tp}(d/A' \cup I')$ and d' realizes p.

But $\mathrm{tp}(b_1, \ldots, b_n/A') = \mathrm{tp}(c_1, \ldots, c_n/A')$ for some finite tuples c_i from $\underline{e}_i \in I$. But then $\varphi(x, c_1, \ldots, c_n)$ is consistent, and we have:

$$\models \forall x(\varphi(x, c_1, \ldots, c_n) \to \psi(x))$$

for all $\psi(x) \in p'(x)$, which contradicts (*). This completes the proof of part (c).

(d) Immediate. $M_{\overline{\kappa}}$ is now known to be a $\overline{\kappa}$-saturated model, of cardinality $\overline{\kappa}$, so can be assumed to be the monster model $\mathfrak{M}$.

(e) Note that a strong type over a set A is precisely a type over $\mathrm{acl}^{\mathrm{eq}}(A)$. We can repeat the proof of (c) working now with types over algebraically closed sets in $\mathfrak{M}^{\mathrm{eq}}$. Alternatively, at least for A the (real) algebraic closure of an infinite subset of $\overline{I}$, A is already an elementary substructure of $\mathfrak{M}$, so complete types over A and strong types over A amount to the same thing.

(f) is just a specialization of (e). $\square$

In particular, for cardinals $\nu \geq \tau$, $\langle M_\lambda : \nu \leq \lambda < \nu^+ \rangle$ is an elementary chain of copies of the saturated model of T of cardinality ν.

Corollary 2.3. *Let T be a complete theory of cardinality τ.*

(a) *Then T is (μ, κ)-almost indiscernible for some $\kappa > \tau$ iff T is (μ, τ^+)-almost indiscernible iff T is (μ, κ')-almost indiscernible for all $\kappa' > \tau$.*

(b) *In particular, under the conditions of a, M_τ is a saturated model which is the algebraic closure of an indiscernible sequence $\langle \underline{e}_\alpha : \alpha < \tau \rangle$ of μ-tuples.*

Comment. We can now assume without loss of generality that T is a complete theory of cardinality τ which is (μ, τ^+)-almost indiscernible for some $\mu \leq \tau$.

Theorem 2.4. *Let T be (μ, τ^+)-almost indiscernible. Then T is stable in every cardinal $\lambda \geq \tau$, hence T is superstable, and in particular if T is countable and $(\aleph_0, \aleph_1)$-almost indiscernible, then T is tt.*

Proof. Once again, we show that T is stable in all cardinals $\lambda \geq \tau$, by the method of the proof of [11, Proposition 2.5].

So let $\lambda \geq \tau$ be a cardinal. Since M_λ is saturated of cardinality λ, it suffices to count the complete types over M_λ. Let $p(v)$ be a complete type over M_λ. Then p is realized in M_{λ^+} by some element d, which is then algebraic over $M_\lambda \cup \{ \underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n} \}$ with $\lambda \leq \alpha_1 < \cdots < \alpha_n < \lambda^+$. So the type of d over $M_\lambda \cup \{ \underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n} \}$ is isolated by some formula $\theta(v, c_1, \ldots c_k)$, where $\theta(v, x_1, \ldots x_k)$ is a formula with parameters from M_λ and $c_1, \ldots c_k$ are entries

from the sequences $\underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n}$. There are at most λ many such formulas θ. By indiscernibility, the type of $\langle \underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n} \rangle$ over M_λ depends only on n. There are only $\mu < \lambda$ choices for finite sequences $c_1, \ldots c_k$ from $\langle \underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n} \rangle$. Hence there are for fixed such θ, no more than μ possibilities for the type of d over $M_\lambda \cup \{ \underline{e}_{\alpha_1}, \ldots \underline{e}_{\alpha_n} \}$, so certainly no more than μ types over M_λ whose realizations are determined by θ. Therefore there are no more than $\mu\lambda = \lambda$ 1-types over M_λ.   □

Comment. For uncountable T, all that follows in general for superstable theories is that T is stable in every cardinal $\geq 2^{|T|}$. So almost indiscernible theories are "strongly" superstable. Any complete theory T in a (possibly uncountable language) which is categorical in $|T|^+$ is "strongly" superstable, cf. the revised edition of Shelah's "Classification Theory", the first paragraph of the proof of [14, Theorem IX.1.15].

We need to make the use of infinitely many variables precise. Let $\vec{v} = \langle v_\beta : \beta < \mu \rangle$ be a sequence of distinct variables. A formula $\varphi(\vec{v})$ "in the variables $\vec{v}$" is some finitary formula φ in some (definite) finite list of variables from $\vec{v}$. This establishes a correspondence between formulas and their variables, and μ-sequences of elements, for the purposes of definitions such as the following:

Definition 2.5. Set $\mathbf{p} = \mathbf{p}(\vec{v}) = \mathrm{tp}(\underline{e}_\omega / M_\omega) = \{ \varphi(\vec{v}) : \models \varphi[\underline{e}_\omega] \}$.

Note that since $\{ \underline{e}_\alpha : \alpha < \overline{\kappa} \}$ is setwise indiscernible, $\langle \underline{e}_\alpha : \omega \leq \alpha < \overline{\kappa} \rangle$ is a Morley sequence over M_ω in $\mathbf{p}$, the so-called *average type* of I_ω.

Proposition 2.6. T *is non-multidimensional.*

Proof. We know that for every ordinal $\lambda \geq \omega$, $\mathrm{tp}(\underline{e}_\lambda / M_\lambda)$ is the nonforking extension of $\mathbf{p}$, and moreover M_λ is saturated (for $\lambda \geq \tau^+$). Let q be over some M_λ, $\lambda \geq \tau^+$. Set $\nu = |\lambda|^+$. As M_ν is ν-saturated, q is realized in M_ν which is in the algebraic closure of M_λ and an independent sequence of realizations of $\mathbf{p}|M_\lambda$ (nonforking extension of $\mathbf{p}$ to M_λ). So q is nonorthogonal to $\mathbf{p}$. This shows that every type is nonorthogonal to M_ω, so T is non-multidimensional.   □

So the number of non-orthogonality classes of regular types is bounded by τ, as T is τ-stable and $|M_\omega| \leq \tau$.

2.2. Structure

We want to prove a version of [11, Proposition 2.10], and explore further consequences of that result. The proposition 2.8 generalizing [11, Proposition 2.8] is essential. In generalizing the proofs of [11, Propositions 2.8, 2.9, 2.10] to our more general setting, we clarify and improve on many steps of the proofs. For background on the "forking calculus" arguments, we refer the reader to [8, Chapter 1, §2], in particular to 2.20–2.29.

Notation. Let $\overline{\mu}$ be $\aleph_0$ if μ is finite and μ^+ if μ is infinite.

Note that $\overline{\mu} \leq \tau^+$.

Recall (cf. [8, Lemma 4.4.1]) that *weight one* types are the ones for which there is a well-defined dimension theory; that non-orthogonality is an equivalence relation on the weight one types, and that in particular every regular type has weight one ([8, Lemma 4.5.3]).

Although the equivalence classes themselves are large (cardinality of the universe), since T is non-multidimensional (Proposition 2.6), we can find a set of representatives of these classes as types over any a-model.

Therefore we can make the following definition:

Definition 2.7. We let $\mathcal{R}$ be the *set* of non-orthogonality equivalence classes of weight one types of T.

If p is some weight one type, then $[p]$ is its class.

Proposition 2.8. *Let $\lambda \geq \omega$ be an ordinal.*

[There are really only two cases of interest, $\lambda = \omega$ and $\lambda = \overline{\mu}$.]

Consider $M_\lambda \prec M_{\lambda+1} = \mathrm{acl}(M_\lambda \cup \underline{e}_\lambda)$.

There is a set of elements $C = \{\, c_j : j \in J \,\} \subset M_{\lambda+1} \backslash M_\lambda$, with J finite if μ is finite and $|J| \leq \mu$ otherwise, such that:

(a) *C is independent over M_λ,*
(b) *each $\mathrm{tp}(c_j/M_\lambda)$ is regular,*
(c) *all regular types occur up to non-orthogonality amongst the various types $\mathrm{tp}(c_j/M_\lambda)$,*

and such that $M_{\lambda+1}$ is a-prime and minimal over $M_\lambda \cup C$.

Without loss of generality, we can fix some set $\mathcal{Q}$ of regular types over M_λ representing the classes of $\mathcal{R}$ over M_λ, and assume that for each $c \in C$, $\mathrm{tp}(c/M_\lambda) \in \mathcal{Q}$, [so that for each $c, c' \in C$, either $\mathrm{tp}(c/M_\lambda) = \mathrm{tp}(c'/M_\lambda)$ or these types are orthogonal].

Proof. Choose $C = \{\, c_j : j \in J \,\} \subset M_{\lambda+1} \backslash M_\lambda$ a maximal independent over M_λ set of elements realizing regular types over M_λ. Note that by Theorem 2.2(f), M_λ is an a-model. Clearly we can make this choice respecting the final statement of the Proposition. $\qquad\square$

Claim 1. *J is finite if μ is finite and of cardinality $\leq \mu$ otherwise.*

Proof. This is a weight argument. In a superstable theory any type of a finite tuple b (over some given base set A) has finite weight in the sense that there is no infinite independent over A set of tuples such that b forks with each of them over A. For if not, forking calculus gives an infinite forking sequence of extensions of $\mathrm{tp}(b/A)$, contradicting superstability. A straightforward extension of this argument shows that if b is a μ-tuple then there is no independent over A set of size μ^+ of tuples each of which forks with b over A. In particular, each c_j, being algebraic over $M_\lambda \cup \underline{e}_\lambda$ forks with $\underline{e}_\lambda$ over M_λ, and so the cardinality of J is at most μ when μ is infinite. $\qquad\square$

Claim 2. *$M_{\lambda+1}$ is a-prime (and therefore prime) and minimal over $M_\lambda \cup \{\, c_j : j \in J \,\}$.*

Proof. Let $N \preceq M_{\lambda+1}$ be the *a*-prime model over $M_\lambda \cup \{\, c_j : j \in J \,\}$. If $N \neq M_{\lambda+1}$ then by Lemma 1.11, there is some $d \in M_{\lambda+1} \backslash N$ whose type over N is regular and does not fork over M_λ, contradicting the maximality of $\{\, c_j : j \in J \,\}$.

If $M_{\lambda+1}$ is not minimal, then $M_\lambda \cup \{\, c_j : j \in J \,\} \subseteq N \prec M_{\lambda+1}$, for some N. But by Proposition 1.10 every elementary extension of M_λ is also an *a*-model, so we can repeat the argument just given to get a contradiction.

For the same reason it follows immediately that $M_{\lambda+1}$ is in fact prime over $M_\lambda \cup C$. $\square$

For the rest, we have already seen (in the proof of Proposition 2.6) that any regular type q is nonorthogonal to $\mathrm{tp}(\underline{e}_\lambda/M_\lambda)$ and so nonorthogonal to a regular type q' over M_λ which is nonorthogonal to p, and so realized in the *a*-prime model over $M_\lambda \cup \underline{e}_\lambda$. But the latter is precisely $M_{\lambda+1}$. So q' is realized in $M_{\lambda+1}$, so forks with $\{\, c_j : j \in J \,\}$ over M_λ. It easily follows that q' is nonorthogonal to some $\mathrm{tp}(c_i/M_\lambda)$. $\square$

When we say that an infinite tuple d is *algebraic* over a set A, we mean that each finite sub-tuple of d is algebraic over A, equivalently that (the range of) d is in the algebraic closure of A.

Proposition 2.9. *Continuing the notation of Proposition 2.8 (with $\lambda = \overline{\mu}$), there are μ-tuples $D = \{\, d_j : j \in J \,\}$ such that:*

(a) *$\mathrm{tp}(d_j/M_{\overline{\mu}})$ has weight one and $c_j \in \mathrm{acl}(M_{\overline{\mu}} \cup \{\, d_j \,\})$ for each $j \in J$;*
(b) *D is $M_{\overline{\mu}}$-independent; and*
(c) *$\underline{e}_{\overline{\mu}}$ is interalgebraic with D over $M_{\overline{\mu}}$.*

[Hence also the types of the tuples d_j represent all the classes of $\mathcal{R}$ over $M_{\overline{\mu}}$.]

Proof. $\underline{e}_{\overline{\mu}} = \langle\, \underline{e}_{\overline{\mu},i} : i < \mu \,\rangle$. Noting that $\mathrm{tp}(\underline{e}_{\overline{\mu}}/M_{\overline{\mu}})$ is the non-forking extension of $\mathbf{p}$ to $M_{\overline{\mu}}$, for the remainder of this proof we will let $\mathbf{p}$ denote this type. Set $C = \{\, c_j : j \in J \,\}$ as given by Proposition 2.8. $\square$

We construct the family D by a sequence of approximations.

Initially choose D so that D is independent over $M_{\overline{\mu}}$ and for each j, d_j realizes $\mathrm{tp}(\underline{e}_{\overline{\mu}}/M_{\overline{\mu}}, c_j)$.

But then D is an independent set of realizations of $\mathbf{p}$ and so $M' = \mathrm{acl}(M_{\overline{\mu}} \cup D)$ is a model: an elementary extension of $M_{\overline{\mu}}$. In particular, as c_j is algebraic over $M_{\overline{\mu}} \cup \{\, \underline{e}_{\overline{\mu}} \,\}$, c_j is algebraic over $M_{\overline{\mu}} \cup \{\, d_j \,\}$. So C is contained in M', and hence $M_{\overline{\mu}+1}$ embeds in M' over $M_{\overline{\mu}} \cup C$. Thus (by taking an automorphism of the universe fixing $M_{\overline{\mu}} \cup C$) we can assume without loss of generality that $M_{\overline{\mu}+1}$ is contained in M'. Hence:

Claim 1. $\underline{e}_{\overline{\mu}} \in \mathrm{acl}(M_{\overline{\mu}} \cup D)$.

We now carry out a construction of parameter sequences f_j which, very informally speaking, encode the domination relation between the pairs c_j and d_j. This will eventually allow us to replace each d_j by a d'_j of weight 1 while preserving all the facts proved so far.

Fix j.

Claim 2. There is a tuple f_j of length at most μ such that f_j is independent from c_j over $M_{\overline{\mu}}$ and c_j dominates d_j over $M_{\overline{\mu}}, f_j$. (That is, if a is independent from c_j over $M_{\overline{\mu}}, f_j$ then a is independent from d_j over $M_{\overline{\mu}}, f_j$).

Proof. This is completely standard. We try to construct a sequence a_α of finite tuples, such that such that for each α, a_α is independent from c_j over $M_{\overline{\mu}} \cup \{a_\beta : \beta < \alpha\}$ but a_α forks with d_j over $M_{\overline{\mu}} \cup \{a_\beta : \beta < \alpha\}$. Notice that then for each α, $\{a_\beta : \beta \le \alpha\}$ is independent from c_j over $M_{\overline{\mu}}$, but d_j forks with a_α over $M_\kappa \cup \{a_\beta : \beta < \alpha\}$. If μ is finite then there is (by superstability) a finite bound on forking sequences of extensions of $\operatorname{tp}(d_j/M_\kappa)$, and in general, one cannot find such a forking sequence of length μ^+. Hence, for some $\alpha < \mu^+$ one cannot continue the construction to get a_α. So take $f_j = \langle a_\beta : \beta < \alpha \rangle$. $\square$

So for each $j \in J$ we can choose f_j as described. But by the choice of D (independent over $M_{\overline{\mu}}$) and the forking calculus we can in fact choose the family of the f_j to be independent from C over $M_{\overline{\mu}}$. Let $\overline{c}$ denote the concatenation of C as a J-tuple, and similarly for $\overline{d}$ and $\overline{f}$. Thus we have:

Claim 3. $\overline{c}$ dominates $\overline{d}, \underline{e}_{\overline{\mu}}$ over $M_{\overline{\mu}}, \overline{f}$, and moreover for each $j \in J$, c_j dominates d_j over $M_{\overline{\mu}, \overline{f}}$.

We can now find (by superstability and considering the cardinality of the relevant set of tuples) a subset A of $M_{\overline{\mu}}$ of cardinality $\le \tau$ (and therefore $< \overline{\mu}$) such that $(\overline{c}, \overline{d}, \underline{e}_{\overline{\mu}}, \overline{f})$ is independent from $M_{\overline{\mu}}$ over A. Let $A' = \operatorname{acl}^{\operatorname{eq}}(A)$, so all types over A' are stationary.

It follows from the basic facts about forking and domination (see [8, Chapter 1, Lemma 4.3.4]) that:

Claim 4. $\underline{e}_{\overline{\mu}} \in \operatorname{acl}(A', \overline{d})$, $\operatorname{tp}(c_j/A')$ is regular, c_j dominates d_j over $A'\overline{f}$, $\overline{c}$ dominates $\overline{d}\underline{e}_{\overline{\mu}}$ over $A'\overline{f}$, and $\overline{c}$ dominates $\overline{c}\underline{e}_{\overline{\mu}}$ over A' (the latter because $\overline{c}$ dominates $\overline{c}\underline{e}_{\overline{\mu}}$ over $M_{\overline{\mu}}$).

Now using the strong $\overline{\mu}$-saturation of $M_{\overline{\mu}}$, let $\overline{f}' \in M_{\overline{\mu}}$ realize $\operatorname{tp}(\overline{f}/A')$. Note that $\overline{c}\underline{e}_{\overline{\mu}}$ was independent from $\overline{f}$ over A', whereby $\operatorname{tp}(\overline{f}'\overline{c}\underline{e}_{\overline{\mu}}/A') = \operatorname{tp}(\overline{f}\overline{c}\underline{e}_{\overline{\mu}}/A')$. Now let $\overline{d}' = (d_j' : j \in J)$ be such that $\operatorname{tp}(\overline{f}'\overline{c}\underline{e}_{\overline{\mu}}\overline{d}'/A') = \operatorname{tp}(\overline{f}\overline{c}\underline{e}_{\overline{\mu}}\overline{d}/A')$. Hence all of Claim 4 holds with $\overline{f}$ replaced by $\overline{f}'$, and $\overline{d}$ replaced by $\overline{d}'$:

Claim 4': $\underline{e}_{\overline{\mu}} \in \operatorname{acl}(A', \overline{d}')$, $\operatorname{tp}(c_i/A')$ is regular, c_j dominates d_j' over $A'\overline{f}'$, $\overline{c}$ dominates $\overline{d}'\underline{e}_{\overline{\mu}}$ over $A'\overline{f}'$, and $\overline{c}$ dominates $\overline{c}\underline{e}_{\overline{\mu}}$ over A'.

Note that as c_j dominates d_j' over $A', \overline{f}'$ we have that

Claim 5. $\operatorname{tp}(d_j'/A'\overline{f}')$ has weight 1.

Now, as $\overline{c}$ is independent from $M_{\overline{\mu}}$ over $A'\overline{f}'$ we have (by domination) that $\overline{d}'$ is independent from $M_{\overline{\mu}}$ over $A'\overline{f}'$, and in particular

Claim 6. $\mathrm{tp}(d'_j/M_{\overline{\mu}})$ has weight 1 for all i (and of course the d'_j are independent over $M_{\overline{\mu}}$).

Finally, by strong $\overline{\mu}$-saturation of $M_{\overline{\mu}+1}$ let $\overline{d}''$ realize $\mathrm{tp}(\overline{d}'/A'\overline{f}'\,\overline{c}\underline{e}_{\overline{\mu}})$ in $M_{\overline{\mu}+1}$. Then the domination statement in Claim 4' implies that $\overline{d}''$ is independent from $M_{\overline{\mu}}$ over $A'\overline{f}'$ and as in Claim 5, each $tp(d''_i/M_{\overline{\mu}})$ has weight 1 and the d''_i are independent over $M_{\overline{\mu}}$. Moreover $\underline{e}_{\overline{\mu}} \in acl(\overline{d}, M_{\overline{\mu}})$ (again by Claim 4'), and as $M_{\overline{\mu}+1} = acl(M_{\overline{\mu}}, \underline{e}_{\overline{\mu}})$ we conclude that $\overline{d}''$ is interalgebraic with $\underline{e}_{\overline{\mu}}$ over $M_{\overline{\mu}}$.

So we replace the family D by $\{\, d''_j : j \in J \,\}$ to conclude the proof of the proposition.

$\square$

Theorem 2.10. *Any model M which contains $M_{\overline{\mu}}$ is the algebraic closure of $M_{\overline{\mu}}$ together with an $M_{\overline{\mu}}$-independent set D of tuples of realizations of weight one types over $M_{\overline{\mu}}$.*

Proof. Fix some model M containing $M_{\overline{\mu}}$. Let C and $\mathcal{Q}$ be chosen over $M_{\overline{\mu}}$ as in Proposition 2.8, for $\overline{\mu}$.

For each $q \in \mathcal{Q}$, let I_q be an enumeration of those $c \in C$ which realize q. So each I_q is nonempty and $M_{\overline{\mu}+1}$ is a-prime and minimal over $M_{\overline{\mu}} \cup \bigcup_q I_q$. Let $\lambda_q \leq \overline{\mu}$ be the cardinality of I_q.

For some sufficiently large λ, let M' be a saturated model of cardinality λ extending M.

Since M extends $M_{\overline{\mu}}$, M is an a-model. For each $q \in \mathcal{Q}$ let I'_q be a maximal $M_{\overline{\mu}}$-independent set of realizations of q in M. Note that it may be empty. Let $C' = \bigcup_{q \in \mathcal{Q}} I'_q$. As in the proof of Proposition 2.8, M is a-prime and minimal over $M_{\overline{\mu}} \cup C'$. Now we can extend each I'_q to I''_q, a maximal $M_{\overline{\mu}}$-independent set of realizations of q in M'. Then I''_q has cardinality λ. Let $C'' = \bigcup_{q \in \mathcal{Q}} I''_q$.

For each $c \in \bigcup_q I''_q$ we can (by Proposition 2.9) choose suitable d_c whose type over $M_{\overline{\mu}}$ is of weight 1 and with c dominating d_c over $M_{\overline{\mu}}$ and $c \in \mathrm{acl}(M_{\overline{\mu}}, d_c)$, such that

(a) if $c \in M$ (i.e. $c \in I'_q$ for some $q \in \mathcal{Q}$) then $d_c \in M$, and
(b) $M' = \mathrm{acl}(M_{\overline{\mu}} \cup \bigcup_{c \in C''} d_c)$.
    Let $B = \mathrm{acl}(M_{\overline{\mu}} \cup \{\, d_c : c \in C' \,\})$. So:
(c) $M_{\overline{\mu}} \cup C' \subseteq B \subseteq M$.

As $\{\, d_c : c \in C'' \,\}$ is independent over $M_{\overline{\mu}}$, it follows that for any choice of $c_1, \ldots, c_r \in C'' \backslash C'$, $\mathrm{tp}(d_{c_1}, \ldots, d_{c_r}/B)$ is finitely satisfiable in $M_{\overline{\mu}}$. It follows that B is the universe of an elementary substructure of M', so by (3) and the minimality of M over $M_{\overline{\mu}} \cup C'$ we have that $M = B$, proving the proposition.

$\square$

Corollary 2.11. *Continuing the notation of the preceding results, we can find (in $M_{\overline{\mu}+1}$) sets $\{\, N(r) : r \in \mathcal{R} \,\}$, each uniquely determined up to isomorphism over $M_{\overline{\mu}}$ by r, with $\mathrm{tp}(N_r/M_{\overline{\mu}}) \in r$ and such that each $N(r)$ is a maximal (with respect to $\subseteq$) weight one set over $M_{\overline{\mu}}$.*

    We call $N(r)$ the hull of r (over $M_{\overline{\mu}}$).

    Furthermore, if M is any model containing $M_{\overline{\mu}}$ (as in 2.10), then M is the algebraic closure of a family (independent over $M_{\overline{\mu}}$) of copies of the various $N(r)$, $r \in \mathcal{R}$.

Proof. For each $r \in R$, let c realize a type in $\mathcal{R}$ over $M_{\overline{\mu}}$, and let $N(r)$ be the a-prime model over $M_{\overline{\mu}} \cup \{\, c \,\}$. The type $q = \mathrm{tp}(N(r)/M_{\overline{\mu}})$ is dominated by $\mathrm{tp}(c/M_{\overline{\mu}})$, so q is weight one and in the class r. If X is any set whose type over $M_{\overline{\mu}}$ is of weight one and in r then the a-prime model over $M_{\overline{\mu}} \cup X$ is isomorphic to $N(r)$, and so $N(r)$ is maximal with the stated properties, and unique up to isomorphism.

    Note that $N(r) \supset M_{\overline{\mu}}$.

    Furthermore, if M is any model containing $M_{\overline{\mu}}$ (as in 2.10), then as in the proof of that theorem we find $M = \mathrm{acl}(M_{\overline{\mu}} \cup \{\, d_c : c \in C' \,\})$ where $\{\, d_c : c \in C' \,\}$ is an $M_{\overline{\mu}}$-independent set of tuples realizing weight one types over $M_{\overline{\mu}}$. Let $r_c = [\mathrm{tp}(d_c/M_{\overline{\mu}})]$. Then for each $c \in C'$, d_c is contained in a copy N_c of $N(r_c)$, which by definition is domination-equivalent to d_c, and so we can choose $\{\, N_c : c \in C' \,\}$ to be independent over $M_{\overline{\mu}}$. Then of course $M = \mathrm{acl}\left(\bigcup_{c \in C'} N_c\right)$. $\qquad\square$

    In particular cases, the structure theory can be refined quite a bit.

Example 2.12. The primordial motivating example for stability theory already exhibits this structure, and more. Let T be the theory of algebraically closed fields of some fixed characteristic. Then a transcendence basis is an indiscernible set, and any model is the algebraic closure of its transcendence basis. So T is $(1, \aleph_1)$-almost indiscernible.

    The results of this section describe a structure theory for the extensions of the model with countably many independent transcendental elements ($\overline{\mu} = \aleph_0$), but of course here we actually have a structure theory for extensions of the prime model.

    $\mathcal{R}$ consists of a unique class, and the hull of that class is the field of transcendence degree one. So Corollary 2.11 sees every algebraically closed field as the algebraic closure of an algebraically independent family of algebraically closed fields of trancendence degree one.

    In the next subsection, we will see similar kinds of examples in theories of modules.

2.3. The case of theories of modules

We take Prest's book [12] as our main reference, to ensure a uniform approach to the subject, with occasional attributions to primary sources.

    Throughout, T is a **complete superstable theory** of R-modules with $|T| = \tau$.

    $\lambda(T)$ is the least cardinal in which T is stable (so $\tau \leq \lambda(T) \leq 2^\tau$).

Proposition 2.13 ([12, Corollary 3.8], due to Ziegler). *If $M \prec N \models T$ then the factor module N/M is totally transcendental.*

    We learn the following facts from Prest [12, §6.5] (originally Pillay–Prest [10]):

Proposition 2.14.   (a) *$M \models T$ is an a-model iff M is pure-injective and weakly saturated.* [12, *Proposition 6.37*]

  (b) *Pure-injective models of T are $\aleph_0$-homogeneous.* [12, *Proposition 6.35*]
  (c) *Elementary extensions of pure-injective models are pure-injective.* [12, *Lemma 6.34*]
  (d) *Elementary extensions of a-models are a-models.* [12, *Corollary 6.41*]

Proposition 2.15. *In general, for any complete theory of modules T', if every a-model is pure-injective, then T' is superstable.* [12, *Proposition 6.40*]

Example 2.16. Consider the p-adics $\overline{\mathbb{Z}_{(p)}}$ in two ways, as a $\mathbb{Z}$-module and as $\overline{\mathbb{Z}_{(p)}}$-module. In both cases the theory is superstable not totally transcendental, with $\lambda(T) = 2^{\aleph_0}$. But the latter has $\lambda(T) = |T|$. (These are used as illustrative examples of many aspects of the model theory of modules throughout [12] and the facts stated here are "common knowledge".)

    Models of the theory as a $\mathbb{Z}$-module have the form $M \oplus \mathbb{Q}^{(\kappa)}$, where $M \preceq \overline{\mathbb{Z}_{(p)}}$ and $\kappa \geq 0$ is a cardinal. Models of the theory as a $\overline{\mathbb{Z}_{(p)}}$-module have the form $\overline{\mathbb{Z}_{(p)}} \oplus \mathbb{Q}_p{}^{(\kappa)}$, where $\mathbb{Q}_p$ is the quotient field.

    In either case, there are no algebraic or definable elements other than 0. Note however that the type of, for instance, $1 \in \overline{\mathbb{Z}_{(p)}}$, while not algebraic, is *limited* in the sense that the pure-injective hull of a realization of it occurs exactly once as a direct summand of any model of T. As a $\mathbb{Z}$-module, the theory is not (μ, κ)-almost indiscernible for any countable μ. But as a $\overline{\mathbb{Z}_{(p)}}$-module, the $2^{\aleph_0}$-saturated model $\overline{\mathbb{Z}_{(p)}} \oplus \mathbb{Q}_p{}^{(2^{\aleph_0})}$ is the *definable* closure of an indiscernible set of tuples of cardinality $2^{\aleph_0}$, where for convenience we take the order type of the tuples to be $2^{\aleph_0} + 2^{\aleph_0}$, the first $2^{\aleph_0}$ components of the tuple being some fixed enumeration of $\overline{\mathbb{Z}_{(p)}}$, and the second sequence of $2^{\aleph_0}$ components ranging over an enumeration of the standard basis for $\mathbb{Q}_p{}^{(2^{\aleph_0})}$. (The details of this construction are made explicit in the proof of Theorem 2.17 following.)

    In fact, the situation described in this example is typical:

Theorem 2.17. *Let T be a superstable theory of modules with $\lambda(T) = |T| = \tau$, (in particular, if T is totally transcendental). Then T is an almost indiscernible theory of modules.*

Proof. We extract the required properties of a-models in superstable theories of arbitrary cardinality from Baldwin's book [1] on stability theory, Chapter XI, §1, §2.

    In particular, there is an a-prime model $\mathcal{A}_0$ (of cardinality τ). There is a saturated proper elementary extension N of $\mathcal{A}_0$ of cardinality τ^+, since T is stable in all cardinals greater than τ. Any type that is realized in $N \backslash \mathcal{A}_0$ is realized by an independent set of cardinality τ^+. The factor module $N/\mathcal{A}_0$ is a tt module by 2.13, which decomposes as a direct sum of indecomposables.

In particular, since N and $\mathcal{A}_0$ are themselves pure-injective, each summand is a direct summand of N. These summands are necessarily of cardinality (less than or equal to) τ, the cardinality of the language, as the theory of $N/\mathcal{A}_0$ is tt, and there are, up to isomorphism, no more than τ distinct summands. There are no "limited" summands (summands which appear a fixed finite number of times in any model) as these all necessarily appear as summands of $\mathcal{A}_0$. Therefore (since N has a large independent set over $\mathcal{A}_0$) every summand of $N/\mathcal{A}_0$ occurs τ^+ times. Let A be the direct sum of one copy, up to isomorphism, of each summand of $N/\mathcal{A}_0$. By the arguments just given, $|A| \leq \tau$. Then $N \cong \mathcal{A}_0 \oplus A^{(\tau^+)}$. Just as we did in Example 2.16, fix an enumeration $\overline{m}$ of $\mathcal{A}_0$ in order type τ and an enumeration $\overline{a}$ of A of order type $\leq \tau$; for $i < \tau^+$ let $\overline{a}_i$ be the copy of $\overline{a}$ on the i-th component of $A^{(\tau^+)}$. Let $\underline{e}_i$ be the concatenation of $\overline{m}$ and $\overline{a}_i$. Clearly $\{\,\underline{e}_i : i < \tau^+\,\}$ is a set of sequences all of the same type and independent, since direct-sum independent, and so is a set of indiscernibles.

Thus T is seen to be (τ, τ^+) almost indiscernible. $\qquad\square$

Comment. In the case where T is tt, we can carry out the construction just described, taking $\mathcal{A}_0$ to be the sum of the limited summands of T, if there are any, or 0 otherwise. So the choice of $\mathcal{A}_0$ as described in the proof of the Theorem does not necessarily give the sharpest possible structure theorem. Nor will this crude construction reveal whether or not T is (μ, τ^+) almost indiscernible for some $\mu < \tau$.

Corollary 2.18. *A complete theory T of modules is almost indiscernible iff it is superstable with $\lambda(T) = |T|$.*

Proof. By Theorems 2.4 and 2.17. $\qquad\square$

Example 2.19. *Consider the following example, used at several places in Prest [12], in particular at Example 2.1/6(vii) (with k a countable field). This was an important example of Zimmermann-Huisgen and Zimmermann [16].*

Let k be an infinite field (of cardinality τ) and set $R = k[(x_i)_{i\in\omega} : x_i x_j = 0$ for all $i, j]$. (We can of course, with some small adjustments to the cardinalities in what follows, make the same construction with an uncountable family of indeterminates.) Then $_RR$ is an indecomposable tt module. Its lattice of pp definable subgroups consists of all the finite dimensional vector subspaces of $J = \bigoplus_{i\in\omega} kx_i$, together with J itself and R. Morley rank equals Lascar rank. Each subspace of dimension n has rank n; J has rank ω and R has rank $\omega + 1$. There are thus two indecomposable pure-injective summands of models of T: $_RR$ corresponding to the types of finite rank/rank $\omega + 1$, and $_Rk$ corresponding to the type of rank ω. In the latter case, the action of R on k is given by $x_i k = 0$ for all i.

The models of the theory of $_RR$ are precisely the modules $_RR^{(\kappa)} \oplus k^{(\lambda)}$, with $\kappa \geq 1$ and $\lambda \geq 0$. For $\kappa \geq \tau$, the saturated model of power κ is $_RR^{(\kappa)} \oplus k^{(\kappa)}$.

So in particular although each free module $_RR^{(\kappa)}$, κ a (non-zero) finite or infinite cardinal, is in the algebraic closure of an indiscernible set of cardinality κ (just take the standard basis vectors), these models are not saturated.

However of course this theory is $(2, \tau^+)$*-almost indiscernible, since each indecomposable is* 1*-generated.*

3. Free algebras

3.1. The general theory

We provide generalizations and extensions of the results in Pillay and Sklinos [11, Section 3] to the uncountable context. None of the proofs depended in any significant way on the assumption that the language is countable, and go through with $\aleph_1$ replaced by τ^+. But we verify all the details in any case. As in [11], we refer the reader to the text [4,3] of Burris and Sankappanavar for the elementary facts of Universal Algebra.

If N is an algebra and $X \subseteq N$, then $\langle\!\langle X \rangle\!\rangle$ is the subalgebra of N generated by X.

We start off with a couple of simple facts about free algebras in a variety. Note that the cardinality of a free basis is *not* in general an invariant of a free algebra unless that cardinality is greater than the cardinality of the language, c.f. Example 3.13.

Lemma 3.1 [11, Remark 3.1]. *Suppose the algebra A is free on $X_1 \cup X_2$, $X_1 \cap X_2 = \emptyset$. Let $A_1 = \langle\!\langle X_1 \rangle\!\rangle$, so A_1 is free on X_1. Let Y_1 be any other free basis for A_1. Then A is free on $Y_1 \cup X_2$.*

We give here a more general version of [11, Lemma 3.7], which is a fact of universal algebra, not a consequence of the context in which we work.

Lemma 3.2. *Let M be a free algebra on τ^+ generators in a variety $\mathcal{V}$ over a language of cardinality τ. Suppose that Y is a free basis of M; $\overline{a}$ is a finite tuple in some other free basis X of M; and for some finite tuple $\overline{y} \in Y$ and finite tuple of terms $\overline{t}$, $\overline{a} = \overline{t}(\overline{y})$.*

Then for any $C \subset Y \backslash \overline{y}$ such that $|Y \backslash C| = \tau^+$, $C \cup \{\overline{a}\}$ may be extended to a basis of M.

Proof. Let $Y_0 = Y \backslash C$. So $|Y_0| = \tau^+$, $\overline{y} \in Y_0$, and $\overline{a} \in \langle\!\langle Y_0 \rangle\!\rangle$. Then $\langle\!\langle Y_0 \rangle\!\rangle \cong M$ as they are both free on τ^+ generators; $\langle\!\langle Y_0 \rangle\!\rangle \subset M$ (and is a proper subset if C is nonempty); and $C \cap \langle\!\langle Y_0 \rangle\!\rangle = \emptyset$ as Y is a free basis. So by Lemma 3.1, $C \cup \{\overline{a}\}$ may be extended to a basis of M. $\square$

So in particular if a belongs to a basis X of M, and Y is a basis of M, then there is $b \in Y$ such that $\{a, b\}$ can be extended to a basis of M.

Context 3.3. Take as a replacement for [11, Assumption 3.2] the following:

Let $\mathcal{V}$ be a variety over an algebraic language $\mathcal{L}$ of cardinality $\tau \geq \aleph_0$. Let the algebra M be a free algebra for $\mathcal{V}$ on a set $I = \{\underline{e}_\alpha : \alpha < \tau^+\}$ (of individual elements), such that M is τ^+ saturated.

Adopt the same notational conventions as in Context 2.1. So in particular the underlying theory T is the theory of M.

Lemma 3.4.   (a) *I is a set of indiscernibles in M.*

(b) *If $I' \subset I$, or if $I' \supset I$ and I' is a set of indiscernibles in $\mathfrak{M}$ extending I, then $\langle\!\langle I' \rangle\!\rangle$ is free on I' in $\mathcal{V}$.*

Proof. (a) holds in general by freeness; the first case of (b) always holds for a free algebra, and the second case of (b) then follows by indiscernibilty and the homogeneity of the universe. For clearly then any subset $I_0 \subset I'$ of cardinality τ^+ is a free basis for $\langle\!\langle I_0 \rangle\!\rangle$; the set of all such subsets of I' is an updirected family; if $f : I' \to A$ for some algebra $A \in \mathcal{V}$ then each $f \restriction I_0$ has a unique lifting to a homorphism $\overline{f_{I_0}} \to A$, and these liftings are pairwise compatible (else we would have a contradiction to indiscernibility). So the union of the maps $\overline{f_{I_0}}$ lifts f to A. $\qquad\square$

Corollary 3.5. (a) T *is superstable.*
(b) *If T is a theory of modules, then T is tt, and $T = \mathrm{Th}(F^{(\aleph_0)})$, where F is the free module in $\mathcal{V}$ on one generator.*

Proof. (a) By Theorem 2.4.

(b) The free module on a set I in $\mathcal{V}$ is isomorphic to $F^{(I)}$, where F is the free module in $\mathcal{V}$ on one generator. So T is a superstable theory of modules closed under products, hence is tt. All infinite weak direct powers of any module are elementarily equivalent, so $F^{(I)} \equiv F^{(\aleph_0)}$. $\qquad\square$

Definition 3.6. Following [11, Definition 3.3], we call $B \subset M$ *basic* if it is a subset of some free basis of M. We call $b \in M$ *basic* if $\{\,b\,\}$ is basic.

Lemma 3.7 [11, Lemma 3.4]. *There is a type p_0 over $\emptyset$ such that for any $a \in M$, a is basic iff a realizes p_0.*

Proof. Since I is an indiscernible set, all the elements of I have the same type p_0. If X is any other basis of M, since $|M| > \tau$, X also has cardinality τ^+, and any bijection between X and I extends, by freeness, to an automorphism of M. So the elements of X also have type p_0.

Conversely, if the type of a is p_0 and $e \in I$ then by saturation there is an automorphism f of M taking e to a; then $f[I]$ is a free basis containing a.
$\qquad\square$

Fix (for the remainder of this section) p_0 as in Lemma 3.7.

Comment. We would like to see that the rank of p_0 is maximal. All that is needed is that if M, N are models of a superstable theory, and $f : M \to N$ is a homomorphism, then the U-rank (Morley rank, as the case may be) of $a \in M$ is greater or equal the U-rank (Morley rank) of $f(a)$. At present, we only have the result for theories of modules.

Lemma 3.8 [11, Lemma 3.8]. *The type p_0 is stationary. Hence so is $p_0^{(n)}$ for any n.*

Proof. We have to show that p_0 determines a unique strong type over $\emptyset$.

Suppose that a and b are realizations of p_0. So a is an element of some basis X of M and b is an element of some basis Y of M. By Lemma 3.2, there is $b' \in Y$ such that $\{\,a,b'\,\}$ is basic, so extends to a basis Z of M. But Z is

indiscernible in M, so a and b' have the same strong type. But b and b' have the same strong type (as elements of Y), so a and b have the same strong type. $\qquad \square$

Corollary 3.9. *All n-types over $\emptyset$ are stationary.*

Proof [From the proof of [11, Corollary 3.10, 3.11]]. If $a \in M$ then $a \in \mathrm{dcl}(e)$ for some finite sequence $e \in I$. But $\mathrm{tp}(e/\emptyset)$ is stationary by Lemma 3.8. Therefore $\mathrm{tp}(a/\emptyset)$ is stationary. $\qquad \square$

Corollary 3.10 [11, Corollary 3.10]. $\mathrm{acl}^{\mathrm{eq}}(\emptyset) = \mathrm{dcl}^{\mathrm{eq}}(\emptyset)$.

Proof. Immediate. $\qquad \square$

Proposition 3.11 [11, Proposition 3.9]. *The sequence I is a Morley sequence in p_0.*

Proof. We have to show that I is independent over $\emptyset$.

Let $I_0 = \{\, \underline{e}_n : n < \omega \,\}$. By homogeneity it is enough to show that $\underline{e}_\omega$ and I_0 are independent over $\emptyset$. Let a realize a non-forking extension of p_0 to I_0. By Lemma 3.2 there is an infinite $I_0' \subseteq I_0$ such that $I_0' \cup a$ is basic. But then by freeness there is an automorphism carrying $I_0' \cup a$ to $I_0' \cup \underline{e}_\omega$, so $\underline{e}_\omega$ and I_0 are independent. $\qquad \square$

Proposition 3.12 [11, Proposition 3.12]. *Let $\bar{a}, \bar{b} \in M$.*

Then $\bar{a}$ is independent from $\bar{b}$ over $\emptyset$ iff there is a basis $A \cup B$, A, B disjoint, of M such that $\bar{a} \in \langle\!\langle A \rangle\!\rangle$ and $\bar{b} \in \langle\!\langle B \rangle\!\rangle$.

Proof. The reverse direction is clear by Proposition 3.11, as any basis is an independent set.

For the forward direction, suppose $\bar{a}$ is independent from $\bar{b}$ over $\emptyset$. Without loss of generality, for some $n < \omega$, $\bar{a}, \bar{b} \in \langle\!\langle \underline{e}_i : i < n \rangle\!\rangle$. In particular, $\bar{a}$ is expressed as a *sequence* of terms in $\{\, \underline{e}_i : i < n \,\}$, $\bar{a} = \vec{t}(\underline{e}_i : i < n)$. Let $\bar{a}' = \vec{t}(\underline{e}_i : n \leq i < 2n)$. Then $\bar{a}'$ is independent from $\bar{b}$ (by the reverse direction already proved!) and $\mathrm{tp}(\bar{a}) = \mathrm{tp}(\bar{a}')$. Thus there is an automorphism of the universe fixing $\bar{b}$ and taking $\bar{a}'$ to $\bar{a}$. The image of the basis I under this automorphism gives us the required decomposition. $\qquad \square$

The proof extends in the obvious way to infinite tuples, and to independence over an arbitrary basic set.

3.2. The particular case of modules

Let $\mathcal{V}$ be a variety of (left) R-modules. The free module $\mathcal{F}_1$ on one generator in $\mathcal{V}$ is clearly an image of the (absolutely) free module on one generator, that is, of ${}_R R$. We take $1 \in R$ as the free generator. So $\mathcal{F}_1 \cong R/I$ for some left ideal I of R. The free module in $\mathcal{V}$ on a set X is then (up to isomorphism) $\mathcal{F}_1^{(X)}$. In particular, by "the free module in $\mathcal{V}$ on κ-many generators", we mean $\mathcal{F}_1^{(\kappa)}$.

Example 3.13. Note however that free modules on different cardinalities need not be distinct: take K a field, κ an infinite cardinal, and set R to be the ring

of all column-finite $\kappa \times \kappa$ "matrices" over K. It is an easy exercise to verify that the usual matrix multiplication is well-defined. By considering partitions of κ into $2, 3, \ldots, n, \ldots$ pairwise disjoint subsets, each of cardinality κ, we see that

$$_R R \cong {}_R R^{(2)} \cong {}_R R^{(3)} \cong \cdots \cong {}_R R^{(n)} \cong \cdots$$

However the cardinality of a free basis is uniquely defined whenever it is infinite.

Let $\kappa \geq |R|^+$. Suppose that $M = \mathcal{F}_1^{(\kappa)}$, the free module on κ generators in $\mathcal{V}$, is saturated and let $T = \mathrm{Th}(M)(= \mathrm{Th}(\mathcal{F}_1^{(\aleph_0)}))$. Since T is a superstable theory of modules and closed under products, it is tt.

Recall the ordering of pp-types of a theory T of modules: if p and q are pp-types of T, then $p \leq q$ if for all $M \models T$, $p[M] \subseteq q[M]$. Equivalently, $p \leq q$ iff $p \supseteq q$. When we say "maximal", we mean "maximal with respect to the ordering $\leq$ on pp-types".

Proposition 3.14. *The pp-type $p_0{}^+$ of a basic element is maximal, hence p_0 has maximal Morley rank.*

Proof. Fix a basis X of M and $e \in X$. So the pp-type of e is $p_0{}^+$. $N \equiv M$ certainly implies $N \in \mathcal{V}$. Let q be a pp-type of T and let $a \in N \models T$ realize q. Define $f : X \to N$ by setting $f(e) = a$ and letting f be arbitrary otherwise. Then by freeness f extends to a homomorphism $\overline{f} : M \to N$, and homomorphisms increase pp types setwise, so $p_0{}^+ \subseteq q$, that is, $p_0{}^+ \geq q$. $\quad\square$

Pillay and Sklinos ask ([11, Question 3.14]) whether the theory of a saturated free algebra must have finite Morley rank, and suggest that the answer should be easy in a variety of R-modules. Under suitable restrictions, the answer is indeed "yes". For instance, if $\mathcal{V}$ is a variety of R-modules such that the free module $_R N_1$ on one generator has a unique indecomposable direct summand, and the free module $_R N$ on $|R|^+$ generators is saturated, then $\mathrm{Th}(N)$ is unidimensional and so has finite Morley rank. However the answer is "no" in general; see Example 3.16 following.

On the other hand it is natural to ask if there is classification of those rings R such that free R-modules on $|R|^+$ generators are saturated. (In fact this was explicitly asked by Piotr Kowalski during a talk on the subject by the second author in Wroclaw, Poland, in July 2019.) This is answered in the theorem below. See Chapter 14 of [12] for definitions of the notions of coherence and perfectness of a ring R. The theorem and proof are thematically close to the material in this Chapter; a partial result in that direction is Exercise 2(a) on page 292.

Recall that the projective modules are precisely the direct summands of free modules.

By the classic theorem of Sabbagh-Eklof [13], cf. Prest [12, Theorem 14.25] (which is stated there for the case of *right* modules), R is left perfect and right coherent iff the class of projective left R-modules is elementary. The underlying algebraic result is due to Chase [5].

Theorem 3.15. *Given a ring R of cardinality τ, $R^{(\tau^+)}$ (regarded as a left R-module) is τ^+-saturated if and only if R is left perfect and right coherent.*

Proof. [$\Rightarrow$] Assume that $_RR^{(\tau^+)}$ is saturated (i.e. τ^+-saturated). By Theorem 2.2 every $R^{(\lambda)}$ for $\lambda \geq \tau^+$ is saturated (and elementarily equivalent to $R^{(\tau^+)}$). Let $T = Th(R^{(\tau^+)})$. As observed after Example 3.13 above T is tt.

We claim that the projective (left) R-modules are precisely the direct summands of models of T. This will suffice, as by [12, Lemma 2.23(a)] this class coincides with the class of pure submodules of models of T, and since T is closed under products, by [12, Lemma 2.31] this latter class is elementary.

If $_RM$ is projective then it is a direct summand of $R^{(\lambda)}$ for some $\lambda \geq \tau^+$, and the latter is a model of T.

On the other hand suppose that M is a direct summand of a model N of T. Take $\lambda \geq \tau^+$ sufficiently large so that N is isomorphic to an elementary substructure of $R^{(\lambda)}$. As T is tt, N is a direct summand of $R^{(\lambda)}$ whereby M is also a direct summand of $R^{(\lambda)}$. So M is projective.

[$\Leftarrow$] Assume that R is left perfect and right coherent (and so the projective left modules form an elementary class). By [12, Corollary 14.22], $T = Th(_RR^{(\tau^+)})$ is tt. As T is tt it has a saturated model M in power τ^+. By freeness, there is a surjection $R^{(\tau^+)} \twoheadrightarrow M$. By the assumption M is projective as it is elementarily equivalent to a free module; so it is isomorphic to a summand, hence to an elementary submodel, of $R^{(\tau^+)}$. As T is tt and nonmultidimensional any elementary extension of a κ-saturated model ($\kappa \geq |T|^+$) is κ-saturated. Hence in particular, $R^{(\tau^+)}$ is saturated. $\square$

Comment. Hence in particular if R is commutative and for some $\kappa > |R| + \aleph_0$, $R^{(\kappa)}$ is saturated, then it has finite Morley rank.

Example 3.16. By contrast, a well-known source of counter-examples (cf. Small [15]) provides us with a saturated free module with infinite Morley rank, and so a counter-example to the question of Pillay and Sklinos [11, Question 3.14]. These examples are almost always presented for *right* modules and we follow that custom here.

Consider the upper triangular matrix ring

$$R = \begin{pmatrix} \mathbb{Q} & \mathbb{Q}(x) \\ 0 & \mathbb{Q}(x) \end{pmatrix}$$

R is well-known to be right artinian and right perfect, but only left coherent, not left noetherian. So by Theorem 3.15, the free (right) module on $\aleph_1$ generators is saturated.

It is a standard exercise to determine *all* the left and right ideals of a ring of this sort.

Furthermore, the pp-definable subgroups of R_R are exactly the finitely generated *left* ideals of R, cf. Prest [12, Theorem 14.16]: if I is generated by $\{r_1, \ldots, r_n\}$, then it is defined in R_R by

$$\varphi(v) = \exists w_1, \ldots w_n \, (v = w_1 r_1 + \cdots + w_n r_n).$$

The Jacobson radical is the "upper right corner" of the ring; it is generated as a right ideal by any non-zero element, but as a left ideal it has the structure of $\mathbb{Q}(x)$ as a vector space over $\mathbb{Q}$. In particular every finite dimensional $\mathbb{Q}$-subspace V of $\mathbb{Q}(x)$ yields a pp-definable subgroup $\begin{pmatrix} 0 & V \\ 0 & 0 \end{pmatrix}$ of the Jacobson radical, and hence we obtain infinite increasing chains of pp-definable subgroups.

Morley rank equals U-rank, cf. Prest [12, Theorem 5.18], and the U-rank of a definable subgroup is just its rank in the lattice of pp-definable subgroups, cf. Prest [12, Theorem 5.12], so the Morley rank of R_R is infinite. [In fact, a more careful analysis shows that the Morley rank is $\omega + 1$.]

4. Questions and open problems

Question 4.1. Is there a fundamental difference between theories that are (τ, τ^+) almost indiscernible and those that are (μ, τ^+) almost indiscernible for some $\mu < \tau$?

In Section 3.1 we get a couple of results giving characterizations of algebraic closure (Corollary 3.10) and independence in a saturated free algebra (Proposition 3.12). The next two questions relate to these results.

Question 4.2. Are there similar results for arbitrary almost indiscernible theories?

Question 4.3. Is there a more general description of independence in the theory of a saturated free algebra?

We are thinking of something that might fit into a general abstract framework similar to that developed for theories of modules in Prest [12, §5.4] (largely based on [9, Pillay–Prest]).

The closest direct analogue of Prest [12, Theorem 5.35] would be the following:

$\overline{a}$ and $\overline{b}$ are independent over $\overline{c}$ iff there is a basis X, the disjoint union of A, B, and C, such that $\overline{c} \in \langle\!\langle C \rangle\!\rangle$ and $\overline{a} \in \langle\!\langle A \cup C \rangle\!\rangle$, $\overline{b} \in \langle\!\langle B \cup C \rangle\!\rangle$.

Question 4.4. Is there any kind of classification of those varieties $\mathcal{V}$ for which the free algebra on τ^+ generators is τ^+-saturated?

One should be cautious, as there are examples in Baldwin-Shelah of such $\mathcal{V}$ which have unstable algebras in the variety.

Question 4.5. What about Question 4.4, assuming the stability of $\mathcal{V}$, that is, that every completion of $\mathrm{Th}(\mathcal{V})$ is stable?

Question 4.6. In Proposition 3.14 we showed that the rank of the type of a basic element of a large saturated free module is maximal. Is this true for large saturated free algebras in general?

For the rest, let us assume that the free algebra M on τ^+-generators is τ^+ saturated, and let $T = \mathrm{Th}(M)$.

Question 4.7. Is T totally transcendental?

Question 4.8. Is there a structure theorem for the algebra M, for example as some kind of a product of a module and of a combinatorial part, along the lines of Hart and Valeriote [6]?

Question 4.9. Implicit in the last few questions is the following:

Is there some kind of relative quantifier elimination theorem for such theories?

Acknowledgements

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

[1] Baldwin, J.T.: Fundamentals of Stability Theory. Perspectives in Mathematical Logic. Springer, Berlin (1988)

[2] Baldwin, J.T., Shelah, S.: The structure of saturated free algebras. Algebra Universalis **17**(2), 191–199 (1983)

[3] Burris, S., Sankappanavar, H.P.: A course in universal algebra. http://www.math.uwaterloo.ca/~snburris/htdocs/ualg.html

[4] Burris, S., Sankappanavar, H.P.: A Course in Universal Algebra, Graduate Texts in Mathematics, vol. 78. Springer, New York (1981)

[5] Chase, S.U.: Direct products of modules. Trans. Am. Math. Soc. **97**, 457–473 (1960)

[6] Hart, B., Valeriote, M.: A structure theorem for strongly abelian varieties with few models. J. Symb. Log. **56**(3), 832–852 (1991)

[7] Pillay, A.: Model theory. Lecture Notes, University of Notre Dame. https://www3.nd.edu/~apillay/pdf/lecturenotes_modeltheory.pdf

[8] Pillay, A.: Geometric Stability Theory, Oxford Logic Guides, vol. 32. The Clarendon Press, Oxford University Press, Oxford Science Publications, New York (1996)

[9] Pillay, A., Prest, M.: Forking and pushouts in modules. Proc. Lond. Math. Soc. (3) **46**(2), 365–384 (1983)

[10] Pillay, A., Prest, M.: Modules and stability theory. Trans. Am. Math. Soc. **300**(2), 641–662 (1987)

[11] Pillay, A., Sklinos, R.: Saturated free algebras revisited. Bull. Symb. Log. **21**(3), 306–318 (2015)

[12] Prest, M.: Model Theory and Modules, London Mathematical Society Lecture Note Series, vol. 130. Cambridge University Press, Cambridge (1988)

[13] Sabbagh, G., Eklof, P.: Definability problems for modules and rings. J. Symb. Log. **36**, 623–649 (1971)

[14] Shelah, S.: Classification Theory and the Number of Nonisomorphic Models, Studies in Logic and the Foundations of Mathematics, vol. 92, 2nd edn. North-Holland Publishing Co., Amsterdam (1990)

[15] Small, L.W.: An example in Noetherian rings. Proc. Nat. Acad. Sci. USA **54**, 1035–1036 (1965)

[16] Zimmermann-Huisgen, B., Zimmermann, W.: Algebraically compact ring and modules. Math. Z. **161**(1), 81–93 (1978)

Thomas G. Kucera
Department of Mathematics
University of Manitoba
Winnipeg MBR3T 2N2
Canada
e-mail: thomas.kucera@umanitoba.ca

Anand Pillay
Department of Mathematics
University of Notre Dame
Notre Dame IN 46556
USA
e-mail: Anand.Pillay.3@nd.edu